

Reading 17: Hypothesis Testing

After completing this reading, you should be able to:

- Construct an appropriate null hypothesis and alternative hypothesis and distinguish between the two.
- Construct and apply confidence intervals for one-sided and two-sided hypothesis tests, and interpret the results of hypothesis tests with a specific level of confidence.
- Differentiate between a one-sided and a two-sided test and identify when to use each test.
- Explain the difference between Type I and Type II errors and how these relate to the size and power of a test.
- Understand how a hypothesis test and a confidence interval are related.
- Explain what the p-value of a hypothesis test measures.
- Interpret the results of hypothesis tests with a specific level of confidence.
- Identify the steps to test a hypothesis about the difference between two population means.
- Explain the problem of multiple testing and how it can bias results.

Hypothesis testing is defined as a process of determining whether a hypothesis is in line with the sample data. Hypothesis testing tries to test whether the observed data of the hypothesis is true. Hypothesis testing starts by stating the null hypothesis and the alternative hypothesis. The null hypothesis is an assumption of the population parameter. On the other hand, the alternative hypothesis states the parameter values (critical values) at which the null hypothesis is rejected. The critical values are determined by the distribution of the test statistic (when the null hypothesis is true) and the size of the test (which gives the size at which we reject the null hypothesis).

Components of the Hypothesis Testing

The elements of the test hypothesis include:

- I. The null hypothesis.
- II. The alternative hypothesis.
- III. The test statistic.
- IV. The size of the hypothesis test and errors
- V. The critical value.
- VI. The decision rule.

The Null hypothesis

As stated earlier, the first stage of the hypothesis test is the statement of the null hypothesis. The null hypothesis is the statement concerning the population parameter values. It brings out the notion that “there is nothing about the data.”

The **null hypothesis**, denoted as H_0 , represents the current state of knowledge about the population parameter that's the subject of the test. In other words, it represents the “status quo.” For example, the U.S Food and Drug Administration may walk into a cooking oil manufacturing plant intending to confirm that each 1 kg oil package has, say, 0.15% cholesterol and not more. The inspectors will formulate a hypothesis like:

H_0 : Each 1 kg package has 0.15% cholesterol.

A test would then be carried out to confirm or reject the null hypothesis.

Other typical statements of H_0 include:

$$H_0 : \mu = \mu_0$$

$$H_0 : \mu \leq \mu_0$$

Where:

μ = true population mean and,

μ_0 = the hypothesized population mean.

The Alternative Hypothesis

The **alternative hypothesis**, denoted H_1 , is a **contradiction** of the null hypothesis. The null hypothesis determines the values of the population parameter at which the null hypothesis is rejected. Thus, rejecting the H_0 makes H_1 valid. We accept the alternative hypothesis when the “status quo” is discredited and found to be untrue.

Using our FDA example above, the alternative hypothesis would be:

H_1 : Each 1 kg package does not have 0.15% cholesterol.

The typical statements of H_0 include:

$$H_0 : \mu \neq \mu_0$$

$$H_0 : \mu > \mu_0$$

Where:

μ = true population mean and,

μ_0 = the hypothesized population mean.

Note that we have stated the alternative hypothesis, which contradicted the above statement of the null hypothesis.

The test statistic

A test statistic is a standardized value computed from sample information when testing hypotheses. It compares the given data with what we would expect under the null hypothesis. Thus, it is a major determinant when deciding whether to reject H_0 , the null hypothesis.

We use the test statistic to gauge the degree of agreement between sample data and the null hypothesis. Analysts use the following formula when calculating the test statistic.

$$\text{Test statistic} = \frac{(\text{Sample statistic} - \text{Hypothesized value})}{(\text{Standard error of the sample statistic})}$$

The test statistic is a random variable that changes from one sample to another. Test statistics assume a variety of distributions. We shall focus on normally distributed test statistics because it is used hypotheses concerning the means, regression coefficients, and other econometric models.

We shall consider the hypothesis test on the mean. Consider a null hypothesis $H_0 : \mu = \mu_0$. Assume that the data used is iid, and asymptotic normally distributed as:

$$\sqrt{n}(\hat{\mu} - \mu) \sim N(0, \sigma^2)$$

Where σ^2 is the variance of the sequence of the iid random variable used. The asymptotic distribution leads to the test statistic:

$$T = \frac{\hat{\mu} - \mu_0}{\sqrt{\frac{\hat{\sigma}^2}{n}}} \sim N(0, 1)$$

Note this is consistent with our initial definition of the test statistic.

The following table gives a brief outline of the various test statistics used regularly, based on the distribution that the data is assumed to follow:

Hypothesis Test	Test Statistic
Z-test	z-statistic
Chi-Square Test	Chi-Square statistic
t-test	t-statistic
ANOVA	F-statistic

We can subdivide the set of values that can be taken by the test statistic into two regions: One is called the non-rejection region, which is consistent with H_0 and the rejection region (critical

region), which is inconsistent with H_0 . If the test statistic has a value found within the critical region, we reject H_0 .

Just like with any other statistic, the distribution of the test statistic must be specified entirely under H_0 when H_0 is true.

The Size of the Hypothesis Test and the Type I and Type II Errors

While using sample statistics to draw conclusions about the parameters of the population as a whole, there is always the possibility that the sample collected does not accurately represent the population. Consequently, statistical tests carried out using such sample data may yield incorrect results that may lead to erroneous rejection (or lack thereof) of the null hypothesis. We have two types of error:

Type I Error

Type I error occurs when we reject a true null hypothesis. For example, a type I error would manifest in the form of rejecting $H_0 = 0$ when it is actually zero.

Type II Error

Type II error occurs when we fail to reject a false null hypothesis. In such a scenario, the test provides insufficient evidence to reject the null hypothesis when it's false.

The level of significance denoted by α represents the probability of making a type I error, i.e., rejecting the null hypothesis when, in fact, it's true. α is the direct opposite of β , which is taken to be the probability of making a type II error within the bounds of statistical testing. The ideal but practically impossible statistical test would be one that **simultaneously** minimizes α and β . We use α to determine critical values that subdivide the distribution into the rejection and the non-rejection regions.

The Critical Value and the Decision Rule

The decision to reject or not to reject the null hypothesis is based on the distribution assumed by the test statistic. This means if the variable involved follows a normal distribution, we use the level of significance (α) of the test to come up with critical values that lie along with the standard normal distribution.

The decision rule is a result of combining the critical value (denoted by C_α), the alternative hypothesis, and the test statistic (T). The decision rule is to whether to reject the null hypothesis in favor of the alternative hypothesis or fail to reject the null hypothesis.

For the t-test, the decision rule is dependent on the alternative hypothesis. When testing the two-side alternative, the decision is to reject the null hypothesis if $|T| > C_\alpha$. That is, reject the null hypothesis if the absolute value of the test statistic is greater than the critical value. When testing on the one-sided, the decision rule, reject the null hypothesis if $T < C_\alpha$ when using a one-sided lower alternative and if $T > C_\alpha$ when using a one-sided upper alternative. When a null hypothesis is rejected at an α significance level, we say that the result is significant at α significance level.

Note that prior to decision making, one must decide whether the test should be one-tailed or two-tailed. The following is a brief summary of the decision rules under different scenarios:


Left One-tailed Test

H_1 : parameter $< X$

Decision rule: Reject H_0 if the test statistic is less than the critical value. Otherwise, **do not reject** H_0 .



Left One-tailed Test




Right One-tailed Test

H_1 : parameter $> X$

Decision rule: Reject H_0 if the test statistic is greater than the critical value. Otherwise, **do not reject** H_0 .



Right One-tailed Test




Two-tailed Test

H_1 : parameter $\neq X$ (not equal to X)

Decision rule: Reject H_0 if the test statistic is greater than the upper critical value or less than the lower critical value.




Two-tailed Test



Consider, $\alpha=5\%$. Consider a one-sided test. The rejection regions are shown below:



Rejection Regions - One-Sided Test



The first graph represents the rejection region when the alternative is one-sided lower. For instance, the hypothesis is stated as:

$$H_0: \mu < \mu_0 \text{ vs. } H_1: \mu > \mu_0.$$

The second graph represents the rejection region when the alternative is a one-sided upper. The null hypothesis, in this case, is stated as:

$$H_0: \mu > \mu_0 \text{ vs. } H_1: \mu < \mu_0.$$

Example: Hypothesis Test on the Mean

Consider the returns from a portfolio $X = (x_1, x_2, \dots, x_n)$ from 1980 through 2020. The approximated mean of the returns is 7.50%, with a standard deviation of 17%. We wish to determine whether the expected value of the return is different from 0 at a 5% significance level.

Solution

We start by stating the two-sided hypothesis test:

$$H_0: \mu = 0 \text{ vs. } H_1: \mu \neq 0$$

The test statistic is:

$$T = \frac{\hat{\mu} - \mu_0}{\sqrt{\frac{\hat{\sigma}^2}{n}}} \sim N(0, 1)$$

In this case, we have,

$$n=40$$

$$\hat{\mu}=0.075$$

$$\mu_0=0$$

$$\hat{\sigma}^2=0.17^2$$


So,

$$T = \frac{0.075 - 0}{\sqrt{\frac{0.17^2}{40}}} \approx 2.79$$

At the significance level, $\alpha = 5\%$, the critical value is ± 1.96 . Since this is a two-sided test, the rejection regions are $(-\infty, -1.96)$ and $(1.96, \infty)$ as shown in the diagram below:



Rejection Regions - Two-Sided Test



Since the test statistic (2.79) is higher than the critical value, then we reject the null hypothesis in favor of the alternative hypothesis.

The example above is an example of a Z-test (which is mostly emphasized in this chapter and immediately follows from the central limit theorem (CLT)). However, we can use the Student's t-distribution if the random variables are iid and normally distributed and that the sample size is small ($n < 30$).

In Student's t-distribution, we used the unbiased estimator of variance. That is:

$$s^2 = \frac{\hat{\mu} - \mu_0}{\sqrt{\frac{s^2}{n}}}$$

Therefore the test statistic for $H_0 = \mu_0$ is given by:

$$T = \frac{\hat{\mu} - \mu_0}{\sqrt{\frac{s^2}{n}}} \sim t_{n-1}$$

The Type II Error and the Test Power

The power of a test is the direct opposite of the level of significance. While the level of relevance gives us the probability of rejecting the null hypothesis when it's, in fact, true, the power of a test gives the probability of correctly discrediting and rejecting the null hypothesis when it is false. In other words, it gives the likelihood of rejecting H_0 when, indeed, it's false. Denoting the probability of type II error by β , the power test is given by:

$$\text{Power of a test} = 1 - \beta$$

The power test measures the likelihood that the false null hypothesis is rejected. It is influenced by the sample size, the length between the hypothesized parameter and the true value, and the size of the test.

Confidence Intervals

A confidence interval can be defined as the range of parameters at which the true parameter can be found at a confidence level. For instance, a 95% confidence interval constitutes the set of parameter values where the null hypothesis cannot be rejected when using a 5% test size. Therefore, a $1 - \alpha$ confidence interval contains the values that cannot be disregarded at a test size of α .

It is important to note that the confidence interval depends on the alternative hypothesis statement in the test. Let us start with the two sided-test alternatives.

$$H_0 : \mu = 0$$

$$H_1 : \mu \neq 0$$

Then the $1 - \alpha$ confidence interval is given by:

$$[\hat{\mu} - C_\alpha \times \frac{\hat{\sigma}}{\sqrt{n}}, \hat{\mu} + C_\alpha \times \frac{\hat{\sigma}}{\sqrt{n}}]$$

C_α is the critical value at α test size.

Example: Calculating Two-Sided Alternative Confidence Intervals

Consider the returns from a portfolio $X = (x_1, x_2, \dots, x_n)$ from 1980 through 2020. The approximated mean of the returns is 7.50%, with a standard deviation of 17%. Calculate the 95% confidence interval for the portfolio return.

The confidence interval is given by:

Then the $1 - \alpha$ confidence interval is given by:

$$\begin{aligned} &= [\hat{\mu} - C_\alpha \times \frac{\hat{\sigma}}{\sqrt{n}}, \hat{\mu} + C_\alpha \times \frac{\hat{\sigma}}{\sqrt{n}}] \\ &= [0.0750 - 1.96 \times \frac{0.17}{\sqrt{40}}, 0.0750 + 1.96 \times \frac{0.17}{\sqrt{40}}] \end{aligned}$$

$$= [0.02232, 0.1277]$$

Thus, the confidence intervals imply any value of the null between 2.23% and 12.77% cannot be rejected against the alternative.

One-Sided Alternative

For the one-sided alternative, the confidence interval is given by either:

$$(-\infty, \hat{\mu} + C_{\alpha} \times \frac{\hat{\sigma}}{\sqrt{n}})$$

for the lower alternative

or,

$$(\hat{\mu} + C_{\alpha} \times \frac{\hat{\sigma}}{\sqrt{n}}, \infty)$$

for the upper alternative.

Example: Calculating the One-Sided Alternative Confidence Interval

Assume that we were conducting one-sided test:

$$H_0 : \mu \leq 0$$

$$H_1 : \mu > 0$$

The 95% confidence interval for the portfolio return is:

$$(\hat{\mu} + C_{\alpha} \times \frac{\hat{\sigma}}{\sqrt{n}}, \infty) = (0.0750 - 1.645 \times \frac{0.17}{\sqrt{40}}, \infty) = [0.03078, \infty]$$

On the other hand, if the hypothesis test was:

$$H_0 : \mu > 0$$

$$H_1 : \mu \leq 0$$

The 95% confidence interval would be:

$$\begin{aligned} &= (-\infty, \hat{\mu} + C_{\alpha} \times \frac{\hat{\sigma}}{\sqrt{n}}) \\ &= (-\infty, 0.0750 + 1.645 \times \frac{0.17}{\sqrt{40}}) = (-\infty, 0.1192) \end{aligned}$$

Note that the critical value decrease from 1.96 to 1.645 due to a change in the direction of the change.

The p-Value

When carrying out a statistical test with a fixed value of the significance level (α), we merely compare the observed test statistic with some critical value. For example, we might “reject H_0 using a 5% test” or “reject H_0 at 1% significance level”. The problem with this ‘classical’ approach is that it does not give us the details about the **strength of the evidence** against the null hypothesis.

Determination of the *p-value* gives statisticians a more informative approach to hypothesis testing. The *p-value* is the lowest level at which we can reject H_0 . This means that the strength of the evidence against H_0 increases as the *p-value* becomes smaller. The test-statistic depends on the alternative.

The p-Value for One-Tailed Test Alternative

For one-tailed tests, the *p-value* is given by the probability that lies below the calculated test statistic for left-tailed tests. Similarly, the likelihood that lies above the test statistic in right-tailed tests gives the *p-value*.

Denoting the test statistic by T , the *p-value* for $H_1 : \mu > 0$ is given by:

$$P(Z > |T|) = 1 - P(Z \leq |T|) = 1 - \Phi(|T|)$$

Conversely, for $H_1 : \mu \leq 0$ the p-value is given by:

$$P(Z \leq |T|) = \Phi(|T|)$$

Where z is a standard normal random variable, the absolute value of T ($|T|$) ensures that the right tail is measured whether T is negative or positive.

The p-Value for Two-Tailed Test Alternative

If the test is two-tailed, this value is given by the sum of the probabilities in the two tails. We start by determining the probability lying below the negative value of the test statistic. Then, we add this to the probability lying above the positive value of the test statistic. That is the p-value for the two-tailed hypothesis test is given by:

$$2[1 - \Phi(|T|)]$$

Example 1: p-Value for One-Sided Alternative

Let θ represent the probability of obtaining a head when a coin is tossed. Suppose we toss the coin 200 times, and heads come up in 85 of the trials. Test the following hypothesis at 5% level of significance.

$$H_0: \theta = 0.5$$

$$H_1: \theta < 0.5$$

Solution

First, note that repeatedly tossing a coin follows a binomial distribution.

Our p-value will be given by $P(X < 85)$ where $X \sim \text{binomial}(200, 0.5)$ with mean 100 ($np = 200 \times 0.5$), assuming H_0 is true.

$$P \left[Z < \frac{85.5 - 100}{\sqrt{50}} \right]$$

$$= P(Z < -2.05) = 1 - 0.97982 = 0.02018$$

Recall that for a binomial distribution, the variance is given by:

$$np(1 - p) = 200(0.5)(1 - 0.5) = 50$$

(We have applied the Central Limit Theorem by taking the binomial distribution as approx. normal)

Since the probability is less than 0.05, H_0 is extremely unlikely, and we actually have strong evidence against H_0 that favors H_1 . Thus, clearly expressing this result, we could say:

“There is very strong evidence against the hypothesis that the coin is fair. We, therefore, conclude that the coin is biased against heads.”

Remember, failure to reject H_0 does not mean it's true. It means there's insufficient evidence to justify rejecting H_0 , given a certain level of significance.

Example 2: p-Value for Two-Sided Alternative

A CFA candidate conducts a statistical test about the mean value of a random variable X .

$$H_0: \mu = \mu_0 \text{ vs. } H_1: \mu \neq \mu_0$$

She obtains a test statistic of 2.2. Given a 5% significance level, determine and interpret the *p-value*


Solution

$$P\text{-value} = 2P(Z > 2.2) = 2[1 - P(Z \leq 2.2)] = 1.39\% \times 2 = 2.78\%$$

(We have multiplied by two since this is a two-tailed test)



Example - Two-Sided Test



Interpretation

The p -value (2.78%) is less than the level of significance (5%). Therefore, we have sufficient evidence to reject H_0 . In fact, the evidence is so strong that we would also reject H_0 at significance levels of 4% and 3%. However, at significance levels of 2% or 1%, we would not reject H_0 since the p -value surpasses these values.

Hypothesis about the Difference between Two Population Means.

It's common for analysts to be interested in establishing whether there exists a significant difference between the means of two different populations. For instance, they might want to know whether the average returns for two subsidiaries of a given company exhibit **significant** differences.

Now, consider a bivariate random variable:

$$W_i = [X_i, Y_i]$$

Assume that the components X_i and Y_i are both iid and are correlated. That is:

$$\text{Corr}(X_i, Y_i) \neq 0$$

Now, suppose that we want to test the hypothesis that:

$$H_0 : \mu_X = \mu_Y$$

$$H_1 : \mu_X \neq \mu_Y$$

In other words, we want to test whether the constituent random variables have equal means.

Note that the hypothesis statement above can be written as:

$$H_0 : \mu_X - \mu_Y = 0$$

$$H_1 : \mu_X - \mu_Y \neq 0$$

To execute this test, consider the variable:

$$Z_i = X_i - Y_i$$

Therefore, considering the above random variable, if the null hypothesis is correct then,

$$E(Z_i) = E(X_i) - E(Y_i) = \mu_X - \mu_Y = 0$$

Intuitively, this can be considered as a standard hypothesis test of

$$H_0: \mu_Z = 0 \text{ vs. } H_1: \mu_Z \neq 0.$$

The test statistic is given by:

$$T = \frac{\hat{\mu}_Z}{\sqrt{\frac{\hat{\sigma}_Z^2}{n}}} \sim N(0, 1)$$

Note that the test statistic formula accounts for the correction between X_i and Y_i . It is easy to see that:

$$V(Z_i) = V(X_i) + V(Y_i) - 2\text{COV}(X_i, Y_i)$$

Which can be denoted as:

$$\hat{\sigma}_z^2 = \hat{\sigma}_X^2 + \hat{\sigma}_Y^2 - 2\sigma_{XY}$$

$$\hat{\mu}_z = \mu_X - \mu_Y$$

And thus the test statistic formula can be written as:

$$T = \frac{\mu_X - \mu_Y}{\sqrt{\frac{\hat{\sigma}_X^2 + \hat{\sigma}_Y^2 - 2\sigma_{XY}}{n}}}$$

This formula indicates that correlation plays a crucial role in determining the magnitude of the test statistic.

Another special case of the test-statistic is when X_i , and Y_i are iid and independent. The test statistic is given by:

$$T = \frac{\mu_X - \mu_Y}{\sqrt{\frac{\hat{\sigma}_X^2}{n_X} + \frac{\hat{\sigma}_Y^2}{n_Y}}}$$

Where n_X and n_Y are the sample sizes of X_i , and Y_i respectively.

Example: Hypothesis Test on Two Means

An investment analyst wants to test whether there is a significant difference between the means of the two portfolios at a 95% level. The first portfolio X consists of 30 government-issued bonds and has a mean of 10% and a standard deviation of 2%. The second portfolio Y consists of 30 private bonds with a mean of 14% and a standard deviation of 3%. The correlation between the two portfolios is 0.7. Calculate the null hypothesis and state whether the null hypothesis is rejected or otherwise.

Solution

The hypothesis statement is given by:

$H_0: \mu_X - \mu_Y = 0$ vs. $H_1: \mu_X - \mu_Y \neq 0$.

Note that this is a two-tailed test. At 95% level, the test size is $\alpha = 5\%$ and thus the critical value $C_\alpha = \pm 1.96$.

Recall that:

$$\text{Cov}(X, Y) = \sigma_{XY} = \rho_{XY} \sigma_X \sigma_Y$$

Where ρ_{XY} is the correlation coefficient between X and Y.

Now the test statistic is given by:

$$\begin{aligned} T &= \frac{\mu_X - \mu_Y}{\sqrt{\frac{\hat{\sigma}_X^2 + \hat{\sigma}_Y^2 - 2\sigma_{XY}}{n}}} = \frac{\mu_X - \mu_Y}{\sqrt{\frac{\hat{\sigma}_X^2 + \hat{\sigma}_Y^2 - 2\rho_{XY}\sigma_X\sigma_Y}{n}}} \\ &= \frac{0.10 - 0.14}{\sqrt{\frac{0.02^2 + 0.03^2 - 2 \times 0.7 \times 0.02 \times 0.03}{30}}} = -10.215 \end{aligned}$$

The test statistic is far much less than -1.96. Therefore the null hypothesis is rejected at a 95% level.

The Problem of Multiple Testing

Multiple testing refers to conducting multiple hypothesis testing on the same data. The reuse of data results in spurious results and unreliable conclusions. Multiple testing results create huge test sizes than the assumed size and thus increasing the probability of type I error and hence the introduction of bias to the test.

Some control methods have been developed to combat multiple testing. These include Bonferroni correction, the False Discovery Rate (FDR), and Familywise Error Rate (FWER).

Practice Question

An experiment was done to find out the number of hours that candidates spend preparing for the FRM part 1 exam. It was discovered that for a sample of **10 students**, the following times were spent:

318, 304, 317, 305, 309, 307, 316, 309, 315, 327

If the sample mean and standard deviation are 312.7 and 7.2 respectively, calculate a symmetrical 95% confidence interval for the mean time a candidate spends preparing for the exam using the t-table.

	q = 0.6	0.75	0.9	0.95	0.975	0.99	0.995	0.9975	0.999	0.9995
n = 1	0.3249	1.0000	3.078	6.314	12.706	31.821	63.657	127.321	318.309	636.619
2	0.2887	0.8165	1.886	2.920	4.303	6.965	9.925	14.089	22.327	31.599
3	0.2767	0.7649	1.638	2.353	3.182	4.541	5.841	7.453	10.215	12.924
4	0.2707	0.7407	1.533	2.132	2.776	3.747	4.604	5.598	7.173	8.610
5	0.2672	0.7267	1.476	2.015	2.571	3.365	4.032	4.773	5.893	6.869
6	0.2648	0.7176	1.440	1.943	2.447	3.143	3.707	4.317	5.208	5.959
7	0.2632	0.7111	1.415	1.895	2.365	2.998	3.499	4.029	4.785	5.408
8	0.2619	0.7064	1.397	1.860	2.306	2.896	3.355	3.833	4.501	5.041
9	0.2610	0.7027	1.383	1.833	2.262	2.821	3.250	3.690	4.297	4.781
10	0.2602	0.6998	1.372	1.812	2.228	2.764	3.169	3.581	4.144	4.587
11	0.2596	0.6974	1.363	1.796	2.201	2.718	3.106	3.497	4.025	4.437
12	0.2590	0.6955	1.356	1.782	2.179	2.681	3.055	3.428	3.930	4.318

A. [307.5, 317.9]

B. [307.6, 317.8]

C. [307.9, 317.5]

D. [307.3, 318.2]

The correct answer is A.

Population variance is unknown; we must use the **t-score**.

To find the value of $t_{\frac{0.05}{2}}$, we use the t-table with $(10 - 1 =) 9$ degrees of freedom and the $(1 - 0.025 =) 0.975$ which gives us 2.262.

So the confidence interval is given by:

$$\begin{aligned} &= \bar{X} \pm t_{\frac{\alpha}{2}} \times \frac{s}{\sqrt{n}} \\ &= 312.7 \pm 2.262 \times \frac{7.2}{\sqrt{10}} = [307.5, 317.9] \end{aligned}$$